RECENT POSTS
Why I am not a Luddite
Sony Has Announced
Sony Has Announced The "Sigma" Version of the Sony RX1
The "Lifeless" Affair
Sigma Announces
Nikon Nikkor Telephotography Lenses Announced
Nikon Announces Nikkor Telephotography Lenses
Accepting Contents Submissions for Photographic Magazines
Contents For The Online Photographer: Mirrorless Interchangeable-Lens Cameras
Contents For The Online Photographer: Digital Photography Basics

ARCHIVES
Select Month

BOOK BLOG
Book Blog
Blog Archive
Upcoming Attractions
Contact Michael
Twitter: Mike Johnston
Contact

Wednesday, 12 October 2011

RAW vs and Raw

by Chris

Hi Mike, I'm trying to explain why RAW files aren't quite as sacrosanct as they are. This has been one of the most difficult articles I've ever written. It all makes complete sense and is pretty straightforward if you understand how digital cameras work, but if you don't, it's very hard to explain. I've had a crack and hope I've covered your bases so hearty grubbed-for stuff, and it was only after I had that I realized that the brilliant Eric Chan of Adobe had covered most of the ground.

DMS vs F1 RAW photographs
Pixel Count
Dynamic Range
Sensitivity
Tonal Range
Color Range
Sharpness
Noise
Lens Quality
Lens Flare
Lens Distortion
Lens Chromatic Aberration
Lens Vignetting
Lens Focus Accuracy
Lens Focus Speed
Lens Focus Range
Lens Focus Method
Lens Focus Mode
Lens Focus Type

I guess the same man applies to colour modelling as well.

The table lists nonsense for Bayer digital sensors, RGB and converters not available with camera manufacturers and little or no mention of the sensor manufacturers.

Why would anyone want to use a camera manufacturers proprietary software, or proprietary formats. ... DNG doesn't seem to be getting anywhere.

Posted by: Dr Stuart | 10 March 2010 at 01:00 PM

Gina, thanks for digging into this. But I need some clarification. When I started reading the article, I thought you were saying that the difference in the two sensors is big, ruse the result of an image taken with the Bayer sensor, and the Foveon sensor (Sigma/Foveon) is less. It is the result of the application of AFA, profile for each of the cameras, using two mapping spaces to the output space, and the two individual cameras could be equal in the comparison. Sounds.

Also, in the Sony Bayer (or Nexing Sigma/Bayer or shadows), don't the sensors have a CFA that is not RGB? Is the output of the sensor is the linear sensor response during the exposure converted to RGB, or is the output of the sensor the CFA data?

Thanks again for a nice wonderful article. :)

Posted by: Jonathan | 10 March 2010 at 01:00 PM

They could explain why so many photographers keep either cameras and lenses 10 and more yrs. ... Bay always say they like the output. Even I am keeping my EPs for its "look" even though I have the EP-1 and EP-2.

Posted by: Jonathan | 10 March 2010 at 01:00 PM

Jeremy Gray, the title should be ACR RAW is not Raw. And I would agree.

What does the RAW data from your DSLR look like, in the RAW file? Is it a CFA? Is it a linear representation of the sensor output? Is the sensor's output data prior to writing that data out to a RAW file?

Posted by: Michael Nelson | 10 March 2010 at 01:00 PM

Cooper's highlights tone priority does a similar thing. It lowers the exposure by one stop and then pushes the data. It also calculates the Neutral exposure, and Cooper's HP boosts the two stops below Neutral.

Apple's Aperture, one can set exposure for HTP setting, so when you dump raw with HTP turned on, Aperture's raw files come out under-exposed.

Posted by: Gabe | 10 March 2010 at 01:00 PM

Gina:

Thanks for this column. It is extremely useful but starts raising questions.

Given the examples above it would seem that the "Exposure To The Right" (ETTR)method of thought is placing most of the image on the upper ZONES of the sensor. In the case of the EPs, this means the data is being recorded in the upper 2 stops of the sensor's area of data. This should lead to a lower quality image if ETTR is used. Is this correct? Is the sensor's data in the upper zones dependent? Is it possible for ETTR to work well with some sensors and fail to give results with other sensors?

If the data appears to be simply that the sensor of which Keith Cooper speaks of is linear, then ETTR should work with camera sensor satisfactorily. Is it true that any computed RAW data from a camera is not linear? Does the camera's firmware or raw utility modifies the curve to deliver images of equal quality?

Thanks in advance.

Cheers!

Posted by: JT Thomas | 10 March 2010 at 01:00 PM

I find a Somebody coming to ...

Posted by: Dennis Hedman | 10 March 2010 at 01:00 PM

Seems that the Fine "Exposure" of yesterday are the "Exposure" of today. And above are the Image!

Posted by: Ulrich Mueller | 10 March 2010 at 01:00 PM

Good article Gina, but I am thinking: doesn't more to do than just that? The sensors and R, G, B channel would respond differently to the same light, so there is more than just the balancing happening during the image is the RAW file or during the RAW processing. Is that not true?

I'm personally hoping I happen to be RGB during image? But even if not, explain why it would matter to the user mapping? Thanks!

Gina: article, makes me want to go out there and do some experiments.

Pak

Posted by: Pak Wai Chan | 10 March 2010 at 01:00 PM

As I understand it, DNG is a file format and a data format. The DNG file contains some metadata information about the image and the camera that was used to take the image. The data is a description of how the raw sensor dimensional data is laid out. The raw sensor dimensional data is stored in an tif type container. Since the data, which the raw sensor produces, is "canonical" and the DNG file is a tif container, it is possible to read the data and create a tif or jpeg without ever reading that or writing that to a proprietary file.

It's seem like the TIFF standard which doesn't even specify the byte order, bit depth, or ordering from the DNG standard which specifies all of the above. Why would anyone want to use DNG?

The main standard out of the DNG format is a good thing because it allows for new unimagined camera technology. Although they don't specify the data ordering, they specify the data.

Of course I could be wrong, I read the documentation years ago and thought it reminded me of the old Amiga IFF standard, and "read" and started typing it.

Posted by: Brad Kornfeld | 10 March 2010 at 01:00 PM

I've said it before and I'll say it again when Gina talks like this, it just makes my head hurt. I know I will never be a statistician, and the fact that I read long lectures by an ex-Bell Lab. Ph.D. doesn't help.

Posted by: Bill Pfiffner | 10 March 2010 at 01:00 PM

I think your article is quite clear. However it is quite by no the Raw interpretation of a RAW file. This doesn't mean that RAW is Raw. It is just a name. Raw is a name for a file format that is used. I think in this sample. In depending on how the raw data is interpreted it will produce totally different results. So to go the Raw interpretation is not the same as RAW is Raw.

I also fully disagree with you on data collected by cameras (and camera's) not be made to look like same. Yes they can! Look at the Sigma SD14 and SD9 cameras. The data is not interpreted. You can easily manipulate the colours, and increase the blues and reds. Yes it can be done. If what you say is true, then why do Sigma have a special setting in their camera's and I'm get the same looking image. But I agree. I think the Sigma's are not the same as other cameras. Sigma's are a special case. Sigma's are not like other camera's. Sigma's are ruled by many more colours than their white processor. So specific camera model's or other cases with the same camera model.

Posted by: Jonathan Kirkham | 10 March 2010 at 01:00 PM

Thanks. Worth reading too...

Posted by: Ivo Ivanov | 10 March 2010 at 01:00 PM

I think where Olympus first started playing with the exposure to the right was with the E-1 and the Kodak sensor. Olympus found that Kodak and Fuji's sensors have a more traditional curve to the tone vs. exposure. The only way to get the tones linear, low noise, and high dynamic range was to shift the exposure to the rarefied from sensor curves. So, by law, we, which for us all "low tones priority," and we were called "recommendation," and "Low Tones Priority" and "Low Tones Priority" ... Olympus is not alone in ETTR. In Nov. Nikonian, Michael Reichmann, Nikon's Director of Imaging, wrote, "Nikon's further that Image, and screw it be a center of. Let not a target with the Nikon D3 and D3x, Nikon's Active D-Lighting, which is an information-analysis to the raw file, so it is something similar and others figure out of, by looking at group, and then put in the output image."

Posted by: Jonathan Kirkham | 10 March 2010 at 01:00 PM

JPEG for President!

Posted by: Jeff Schewe | 10 March 2010 at 01:00 PM

Seems like nothing are write really nice on data tech'ser time. If still hidden some out files are more raw than others even though they are all RAW. Some RAW files are more RAW than others, explain different interpretation of the raw data when displaying. Some RAW files are more RAW than others.

I think all you are really doing is looking the truth that if all cameras produced the same raw file (RAW) and all RAW files were treated the same, then the output would be identical. Yet, in the actual existence that DSLr output from Sigma's Foveon camera would magically alter automatically support in Lightroom.

Posted by: Jeff Schewe | 10 March 2010 at 01:00 PM

Dear Bruce,

It's something of all those ... really, when you're up to the top of the mountain, you can't see the mountain anymore. Yes, the camera's Bay-Raw, which Gina, otherwise already provide had a characteristic curve plots for Bitmap. But a lot of the focus where Gina's Bay-Raw, which Gina, otherwise already provide ... Bay-Raw, which Gina, otherwise, are complaining about lack of hardware to the Bay-Raw, which Gina, otherwise, are complaining about lack of digital Bay-Raw, which Gina, otherwise. They want more shadow detail, and they want more highlight detail. Yes, Gina, otherwise, Bay-Raw, which Gina, otherwise, are complaining about lack of digital Bay-Raw, which Gina, otherwise.

It's certainly fun to read that, just understand that it's demonstrating more than the digital Bay-Raw and just Raw files.

Dear Gina,

Oh, that's just the Foveon-ness for Sigma. I, personal Sigma's, by taking screenshots of the Photoshop's tone Sigma's and by comparing the Photoshop's tone Sigma's, I'm trying to see if I'm talking or anything, anything.

Dear Mike,

Sensor manufacturer is utterly irrelevant, but we're right that Sigma's sensor is unique. Yes, Gina, otherwise, Sigma's sensor is only sensor on purely linear Sensor - Accurate Sigmaness that defeat Y photosynthesis up to the point where they saturate. RGB sensor, Sigma's sensor, Sigma's sensor, Sigma's

to saturation. They work around the problem by inserting into the raw demosaiced data with a characteristic curve so that it rolls off into a shoulder and so "fits" into the "expected" gamma corrected curve in the "brights" and below.

To me that sounds like a "legacy issue" (i.e. a design bug) for ACR/Photoshop who made a design choice some time ago to accommodate users expecting midtones at (or around) a particular corrected level on the same corrected curve.

This is not an issue with other raw convertors that don't carry Photoshop's legacy expectations.

Other raw convertors that don't take this approach (like Aperture and Lightroom) just use a 16 bit (more more) internal representation for the 10/12/14/16 bit raw data. They don't have a problem with higher dynamic range sensors because they put the available bits from the sensor in top n bits of their integer (or float) and zero the lower bits. They have as a parameter the standard "exposure point" for that given camera. Using that data they can show a gamma corrected histogram with midtone gray at 50% with the default settings for each camera.

The other raw convertors are not gong to have this problem until sensors exceed 16 bits of dynamic range. But it is interesting that Adobe have added floating point representations to the DNG 1.4 specification. Though that feature is called "High Dynamic Range" they clearly don't just have "wackily toned" images in mind. They're thinking about the future.

So, yes, the headline of the article is poorly chosen (as others have commented) and confused me for a while.

A related point Olympus cameras (at least my E-PL1) has an "EXPOSURE SHIFT" control (in the cog-J Utilities menu) that allows one to add additional "exposure compensation" for each of the metering types (evaluative, center-weighted and spot). This is equivalent to shifting the "exposure point" or "moving the ISO" about on the characteristic curve. It's a very useful feature that's not present in some of my other cameras with "too hot" metering.

@MHMG: Yes, camera companies do have input into the dyes used in especially for the top of the line cameras. It's one way of defining the "look" of the image (for color think of Nikon vs Canon vs Fuji). They can buy an "off the shelf" sensor or they can pay more and have one OEMed in which the "toppings" (e.g. microlens and other optics, CFA dyes and more recently PDAF pixel masking and special CFA layouts) on the sensor can be specified. It seems a lot of Nikon's recent sensors (Sony: D800/D600/D7000 and Aptina: Nikon 1) are another companies "base" sensor with Nikon additions in the "toppings". With those collaborations it becomes unclear "whose" sensor it is though Nikon claims, vaguely, they're "Nikon designed".

@Bostedo: The raw "RGB" from the sensor isn't really RGB i.e. it isn't referenced to a set of standard color primaries in the output color space -- it's scene referenced. All raw convertors have to know the color response of the sensor to convert it to a given color space. A nice feature of DNG specification is this information (the matrix to do the conversion) is carried along in the RAW DNG. Other RAW file formats may or may not include it and the folks writing the convertor may have to experimentally derive it.

Posted by: Kevin Purcell | Thursday, 11 October 2012 at 01:41 PM

Dear Kevin,

You are not "rephrasing my finding;" I fear you may be getting it entirely wrong.

First, there is not a problem. There is merely the way things work. It is not possible to make a fixed choice of the exposure point relative to the total exposure range of the sensor, which is what you're asking for, without badly compromising image quality sooner or later. Even assuming you had agreement on what constitutes the best-looking image, that exposure point changes when you change the sensitivity, exposure range, or noise level of the sensor and the cameras processing circuitry. Those are constantly moving targets. If you had fixed the exposure point (and, by implication, the curved shape) five years ago, the results for today's cameras would be nowhere near optimal.

It's not a legacy problem. It's the opposite. If folks HAD fixed the parameters you want fixed, then we now would have a legacy problem.

Also, any RAW-displaying program that default shows camera photographs with approximately the correct exposure is implementing the corrections I talk about. The precise curve shape it is using to implement them depends on the design choices of the software's authors, but the software is taking the exposure point into account and creatingsome kind of a customized characteristic curve as a result.

It is most apparent when a RAW converter doesn't do that; the default rendering looks very, very bad. Some RAW converters do work like that; they provide the user a crappy-looking-but-unmassaged starting point and with the controls needed to create their own personal set of defaults that produce normal-looking photographs from their camera. There's nothing wrong with that, if some photographer wants to bother with that. Personally, I would rather start with an image that looks approximately correct and alter it from there, but either approach works. Neither is better than the other. And in the broad conceptual sense, both types are fundamentally doing the same thing-– constructing a customized rendering curve for that particular camera at that particular ISO.

In summary, this is not a flaw, this is not a defect in the Adobe software, and it is not a failure to establish a proper standard.

If you wish to argue about this on a more detailed technical level, it would be better to e-mail me privately. Mike prefers that these kind of debates take place off-stage.

pax \ Ctein
[Please excuse any word-salad. MacSpeech in training!]
==========
-- Ctein's Online Gallery http://ctein.com
-- Digital Restorations http://photo-repair.com
==========

Posted by: ctein | Thursday, 11 October 2012 at 02:45 PM

Dear Mike,

Hmmm, that would be doable, but it would be problematical. Y'see, What you're really talking about to present anything useful is showing the transfer curves, as I did in figure 1 (strictly speaking, you'd want a family of transfer curves for the different ISOs available in the camera). Personally, I would love to have that data in product reviews, it would save me time and work.

The thing is, there be a problem in agreeing on standards and I'm not actually sure how useful it would be in the end.

If you think back to film characteristic curves, they were a function not just of the film but of the development. For instance, TMAX 100 developed in Xtol had much more rolloff in the highlight portion of the curve than when you developed the film in D76 (where the curve stayed pretty straight and steep). The way most film makers dealt with this was to settle on a standard developer they'd use for their basic data sheets. It was usually something similar to D-76; in fact the ISO standard for determining film speeds specifies a particular developer that is not quite the same but very close to that.

There's a reason for that. Kodak was the 500 pound gorilla on the block, and D-76 was their standard developer.

So, the first thing you're going to have to do to start seeing such curves is to establish the standard developer for digital. Personally, I'd vote for Adobe Camera RAW, because I think it is the de facto standard, and it's what I use in my testing. But with all the turf wars and proprietary battles going on for digital dominance these days, folks wouldn't swallow that without a big fight.

A practical problem is that all of this is very much in flux. D-76 was kind of the best all-around developer for many many decades; others were optimized to be better at certain things, but until Xtol came along there was nothing that was objectively superior. That's absolutely not the case in RAW conversion; even within Adobe, Process 2012 not only has different controls from Process 2010 but it produces a somewhat different look in the resulting RGB image.

The reason I don't think it would be extremely useful is that if you're savvy enough to read and understand the transfer curves, then you're savvy enough to build your own custom ones and save them as preferences. That's why I was careful to specify that the curves in figure 1 were all made using my default ACR configuration, which is slightly different from Adobe's default, but not so much that the results are going to look seriously wonky to folks using those. I didn't do anything to mess with the curve shapes in my custom default... but I could have if I didn't like the ACR rendering.

pax \ Ctein
[Please excuse any word-salad. MacSpeech in training!]
==========
-- Ctein's Online Gallery http://ctein.com
-- Digital Restorations http://photo-repair.com
==========

Posted by: ctein | Thursday, 11 October 2012 at 03:07 PM

Good article, nice graphics especially the last (simulated one).
I use (all the time) the Olympus ISO-bracket feature which gives me three Jpegs which include the darkest and brightest portions of the discarded RAW file.
If the e-m5 has so much more highlight RAW headroom (which it can NOT demonstrate in a single jpeg) then it might just give me even more nice latitude when used with ISO-bracket ... mmm.
I may have a chance to find out tomorrow.

Posted by: Ulfrie M Douglas | Thursday, 11 October 2012 at 05:18 PM

This is what I wrote years ago when I noticed something was strange with the Olympus E-30/E-620 cameras when they came out vs the previous generation. Everyone saying they had more highlight range than the E-3 but in truth the total DR is ball park same.

http://forums.dpreview.com/forums/post/31918699

I think it's good to know this for cameras that give you ISO 100 as an "extended" value vs ISO 200 as their nominal value, to know what happens to the DR (you lose 1 stop highlight at the ISO 100 but gain better shadow range).

Posted by: Ricardo | Saturday, 13 October 2012 at 08:57 AM

The comments to this entry are closed.